Exhibit J

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA

UNITED STATES OF AMERICA, ET AL., **Plaintiffs,** **v.** **GOOGLE LLC,** **Defendant.**	Civil Action No.: 1:23-cv-0108 (LMB/JFA)

EXPERT REPORT OF PAUL R. MILGROM

DATE: January 23, 2024

HIGHLY CONFIDENTIAL
SUBJECT TO PROTECTIVE ORDER

3. Plaintiffs' Theory of Harm to Publishers and Competing Exchanges Is Based on Faulty Assumptions

237. Plaintiffs' experts claim that Poirot "reduc[ed] scale for rival exchanges," and that this reduction in scale "constricted publishers' choices for ad exchange services."[453] Plaintiffs also point to pre-launch experiments projecting that "Project Poirot would reduce publisher display revenue from DV360 by over 10%," as part of a proposed theory of harm to publishers and competing exchanges.[454]

238. This theory of harm is based on the Plaintiffs' persistent (but faulty) assumption that advertisers do not change their bids in response to changing incentives. Poirot was a free service to DV360 advertisers and, in its absence, advertisers would be incentivized to pursue their own bid optimization programs and these would most likely be less efficient than the one that DV360 designed. Such self-service is the relevant comparison to assess the effects of Poirot, not the Plaintiffs' fictitious but-for world in which advertisers do not respond to incentives.

239. In the absence of Poirot, advertisers that continued to report the same fixed CPMs to DV360 despite exchanges transitioning away from second-price auctions would experience large reductions in their profits from online display advertising. For example, an advertiser that used DV360 to bid into an exchange that suddenly transitioned from using a second-price auction to using a first-price auction would notice—in the absence of Poirot—that it suddenly started paying more for impressions than it did before: in particular, a price equal to its bid. If that bid was optimized for a second-price auction, so

[453] Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 237.

[454] Complaint ¶ 221.

that it equaled its value for the impression, the advertiser would suddenly find itself earning *zero* advertiser surplus on the impressions it won. Clearly, that advertiser would be incentivized to respond, either by excluding that exchange or by reducing the fixed CPM it reports to DV360 to use for bidding. As described earlier, Plaintiffs' experts acknowledge advertisers' incentives to shade their bids into such an exchange.[455]

240. In the absence of a program like Poirot, each advertiser would face the complex task of identifying optimal bids on its own, which would require costly experimentation and engineering resources. Such experimentation was made more complicated by the presence of exchanges using dirty auctions, which sought to obscure their auction rules. Even so, at around the time Poirot was introduced, other buying tools had already started to develop bid optimization programs for non-second price exchanges.[456] Many DSPs including [redacted] implemented programs similar to Poirot, indicating that bid optimization is perceived as a valuable service for advertisers.[457]

[455] *See* Expert Report of R. Abrantes-Metz (Dec. 22, 2023), at ¶ 281 ("It is not optimal for advertisers to bid their willingness to pay in first-price auctions, or those that resemble first-price auctions, because advertisers will not extract any surplus from the auction. Consequently, it can be efficient for buyers to 'shade' their bids below their willingness to pay when bidding into first-price auctions."); Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 220 ("By the mid-2010s, Google believed that some exchanges had moved away from running 'clean' second price auctions, and that some were running non-transparent combinations of first- and second price auctions. Recall that in first price auctions, it is optimal for bidders to reduce their bids[…] "); Expert Report of R. Lee (Dec. 22, 2023), at ¶ 240 ("In an auction with both a soft- and hard-floor, a bidder's bid thus affects not only the likelihood of winning but also, in some cases, the auction price. As a result, always bidding one's own valuation in an auction with a soft floor is no longer optimal and some bidders can do better by shading their bids.").

[456] *See* Sarah Sluis, Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price, AdExchanger (Sep. 05, 2017), https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/ ("To combat price increases, some buyers have already started shading, or reducing bid prices. But that strategy comes with its own risks: Buyers will lose out on inventory they want if they submit too low a bid and the auction turns out to work with second-price logic.").

[457] [redacted]

241. Poirot made bidding easier for DV360 advertisers by performing bid optimizations for them, and more efficiently. DV360 and other large DSPs could take advantage of the scale economies involved in experimentation to adapt bids to more information than any single advertiser observes, allowing them to bid more efficiently on behalf of their advertisers than any individual advertiser could do on its own. In this way, Poirot (and similar programs used by other buying DSPs) reduced the costs to advertisers of optimizing their bids and increased efficiency by reducing the likelihood of bidding errors.

242. To assess any harms of Poirot to publishers and other exchanges, a more revealing question is whether Poirot caused the payments to exchanges using non-second-price auction rules to fall, on average, below the payments for similar impressions on a second-price exchange. Auction theory suggests that the answer is "no." In the standard independent private values auction model used by the Plaintiffs' experts,[458] the celebrated

[458] *See* Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 56 ("An important assumption in the auctions literature, adopted by most of the literature in display advertising in particular, is that buyers' valuations are independent and private. This means that each buyer's valuation is privately known to themselves at the time of the auction and is independent of competitors' (*i.e.*, other competing buyers') valuations."). *See also* Expert Report of R. Lee (Dec. 22, 2023), at ft. 316 ("The type of auctions described in this section and analyzed in this report is known in economics as independent private value auctions, where each bidder's valuation is privately known and not affected by information possessed by other bidders. Auctions for digital advertising have often been modeled by economists and researchers as independent private value auctions.").

Revenue Equivalence Theorem implies that the average revenue for publishers and exchanges is the same with profit-maximizing bids (which Poirot was designed to achieve) in a first-price auction as in a second-price auction.[459] Poirot helps advertisers to bid optimally but would not be expected to reduce other exchanges' average prices below the same market-clearing level that is achieved by second-price auctions.

243. I present an especially simple example to illustrate the power of this *revenue equivalence* idea. Suppose that there are two bidders: Bidder 1 with a value of $1.00 CPM for each impression, and Bidder 2 with a value of $2.00. Each bidder knows its own value but not the other's value. Suppose the publisher floor price is less than $1.00.

244. In a second-price auction, both bidders behave optimally by bidding their values. The seller is paid the value of the second-highest buyer, which is $1.00. In a first-price auction, both bidders could experiment to learn how to bid optimally. Bidder 1 will never experiment by bidding higher than its value of $1.00, because any higher bid can never yield a positive surplus. Bidder 2's experiments will reveal that it always wins the impression when it bids at least $1.00, so it will learn not to bid more than that. It may also experiment with bids less than $1.00 and find that those bids too often lose the auction. This process of experimentation would ultimately result in Bidder 2 learning to bid just above $1.00 to always win the auctions. The resulting per-impression publisher revenue in the first-price auction will be around $1.00, the clearing price in the second-price auction.

[459] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, pp. 73-77. *See also* Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58-73; Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, *13*(3), 227-286; McAfee, R. P., & McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, *25*(2), 699-738.

245. For both kinds of auctions, the advertiser with the $2 value eventually learns to bid optimally, winning most of the impressions and paying about $1. Although my example is very simple—it includes only first- and second-price auctions and makes learning easy with its assumption that values do not change from impression to impression—it illustrates a general principle: bid adjustments by optimizing bidders may fully offset the direct effects of changes to auction rules. Omitting bidders' attempts to optimize can mislead, so a full analysis must account for incentives to do so.

246. Finally, Plaintiffs in their Complaint (but not their experts) argue that Poirot reduced the quantity of impressions sold, claiming "Poirot prevented these advertisers from spending their full advertising budgets and resulted in some ad inventory going unfilled—a loss for the display advertising market as a whole."[460] This is wrong. It neglects the fact that, when Poirot helped advertisers avoid overpaying for impressions, the spending saved on impressions purchased at a lower price due to Poirot could be spent on other impressions. Google engineers expected this would occur, writing, "we expect the budget server to react quickly to adjust impression probabilities and bring [budget-constrained advertisers] back up to spending their budgets."[461] One of Plaintiffs' experts, Professor Weintraub, also acknowledges this expectation.[462] Plaintiffs and their experts have not presented any

[460] Complaint ¶ 221.

[461] Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -814. To understand how this effect could occur, imagine that, without Poirot, an advertiser had a $100 budget and won 1000 impressions, for an average cost per impression of $0.10. Optimal bidding under Poirot into an exchange using non-second-price auctions might allow the advertiser to win the same 1000 impressions for a lower average cost per impression of, say, $0.09, leaving $10 left in the advertiser's budget, which could be used to purchase an additional 111 impressions at that same average cost.

[462] *See* Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 233 ("Because Poirot reduces DV360's bids on some of [*sic*] non-AdX exchanges, DV360 reduced its ad spend on these non-AdX exchanges. As a result, some advertisers were left with unspent budgets. Many advertisers using DV360, however, tasked it to spend all of a budget they set. Google made budget adjustments to the experimental results that would account for advertisers exhausting budget targets on DV360.").

contrary evidence, and this adjustment would result in *more* impressions being filled, not fewer.

4. Poirot Applied Equally to AdX, as Well as to Other Exchanges

247. Plaintiffs' experts argue that Poirot was discriminatorily applied because it did not modify bids into AdX, even though AdX "did not implement clean second-price auctions."[463] They further claim that Google's sell-side and buy-side teams worked together to prevent Poirot from detecting any auction distortions on AdX.[464] To the contrary, from September 2017 onwards (just two months after the program's launch), Poirot applied the same experiment-and-optimize strategy to AdX as to any other exchange. Using the 10% profit threshold, it found that bidding truthfully was an optimal or near-optimal bidding strategy on AdX, and accordingly did not shade its AdX bids. Similarly, Poirot did not adjust bids into non-Google exchanges such as United and Improve Digital, where optimal bid shading was also predicted to increase advertiser surplus by less than 10%.[465] Indeed, as Plaintiffs' experts note,[466] Poirot acted as a limit

[463] *See* Expert Report of R. Ravi (Dec. 22, 2023), at III.D.4 ("AdX Did Not Run Clean Second-Price Auctions Yet Google Did Not Decrease Bids on AdX"); Expert Report of R. Abrantes-Metz (Dec. 22, 2023), at ¶¶ 284-285; Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 211.

[464] *See* Expert Report of R. Ravi (Dec. 22, 2023), at ¶ 206 ("Google's internal documents, as well as the source code, indicate that Google kept AdX's DRS feature turned off on DV360[] such that it would not treat AdX as an exchange running dirty second-price auctions."), ¶ 207 ("Google tweaked its products' functionalities to avoid potential conflicts between its exchange and buy-side tools and gave preferential treatment when needed […] [This] directly contravenes Google's argument that Poirot was not active on AdX due to Google running clean second-price auctions.").

[465] See typical experiment on these exchanges displayed in Figure 8.

[466] Expert Report of G. Weintraub (Dec. 22, 2023), at ¶ 226 ("In its internal discussions, Google engineers contemplated changing Poirot but ended up proposing a 'watered down' version of the program that would make RPO undetectable to Poirot.").

on AdX's designs, with AdX deciding to reject a version of its Reserve Price Optimization program (RPO) in order to avoid triggering "things like Poirot."[467]

248. Figure 8 displays results from "typical" Poirot v1 experiments, which show the gains from bid shading on different exchanges.[468] While DV360 could add more than 50% to advertiser surplus by shading bids into [redacted] and [redacted], shading bids into AdX and Improve Digital would add less than 10%. Subsequent versions of Poirot made similar findings.[469]

[467] *See* Meeting Notes, "AdX Dynamic Price Meeting Notes" (Apr. 11, 2018), GOOG-AT-MDL-012701069, at -073 ("Online RPO […] Three potential less aggressive versions […] Waiting on gTrade team to let us know which versions are acceptable for Poirot"), -075 ("Online RPO […] Working on less aggressive version to avoid things like Poirot").

[468] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -641.

[469] *See* Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -683 ("The new poirot [sic] model turns out to be a no-op on AdX similar to the prod model.").

b. tDRS predicts that there will not be an AdX buyer who can beat the price $r/0.64$, and applies a 0% per-impression revenue share. In this case, the AdX buyer wins the impression if it bids above the floor price of $r/0.8$, exactly the same as the pre-revenue share for advertisers before DRS. Because tDRS incentivizes truthful bidding, the AdX buyer wins the impression if and only if it would win it without tDRS. In this case, the payment to the publisher for the impression is $r/0.8$ and the publisher accrues a "debt" of $0.2\ r/0.8$, leading to a net payoff of r. The expected publisher revenue is therefore unchanged by tDRS.

569. Thus, the impression sells with the same probability under tDRS as without DRS, and the publisher earns more revenue than without DRS.[900] Under the optimal floor price, which may differ from $r/0.8$, the publisher can only receive even higher revenues. ■

570. Although an assumption of Theorem 9 is that AdX is perfectly able to predict whether its buyers will clear a given floor price, the same conclusion holds as long as (1) AdX predicts bids better than publishers and (2) the revenue share under tDRS maximizes Google's expected profits (or, equivalently, the expected revenues of the publishers).

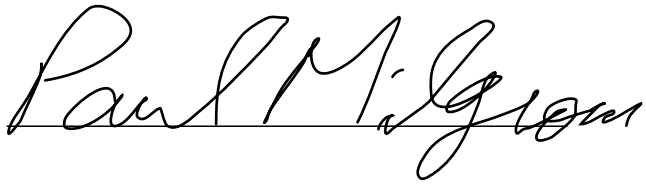

Paul R. Milgrom, Ph.D.

January 23, 2024

[900] In this analysis, I have fixed the bids of header bidders after the introduction of tDRS. In general, header bidders need to account for the change of format in the AdX auction because it may affect their expected payoff of any given bid. However, because under the posited strategy of reporting floor r/0.8, there is no change in the probability of sale on AdX, the header bidders' bids should not change either. Under the publisher's optimal strategy, however, header bidding advertisers' bids would change, but this does not affect the conclusion of Theorem 9.